arXiv:1609.09435v1 [cs.SI] 29 Sep 2016

On the statistical properties of viral misinformation in online social media

Alessandro Bessi*

[a] University of Southern California, Information Sciences Institute, Marina del Rey, Los Angeles, CA, USA
[b] IUSS Institute for Advanced Study, Pavia, ITALY
[c] IMT Institute for Advanced Studies, Lucca, ITALY

Abstract

The massive diffusion of online social media allows for the rapid and uncontrolled spreading of conspiracy theories, hoaxes, unsubstantiated claims, and false news. Such an impressive amount of misinformation can influence policy preferences and encourage behaviors strongly divergent from recommended practices. In this paper, we study the statistical properties of viral misinformation in online social media. By means of methods belonging to Extreme Value Theory, we show that the number of extremely viral posts over time follows a homogeneous Poisson process, and that the interarrival times between such posts are independent and identically distributed, following an exponential distribution. Moreover, we characterize the uncertainty around the rate parameter of the Poisson process through Bayesian methods. Finally, we are able to derive the predictive posterior probability distribution of the number of posts exceeding a certain threshold of shares over a finite interval of time.

Keywords: misinformation, online social media, extreme value theory

1. Introduction

The wide availability of user-provided contents in online social media encourages the aggregation of people around common interests and narratives. The direct path from producers to consumers of contents drives the emergence of a disintermediated enviroment that is changing the way people become informed, interpret facts, and form their opinions [1, 2, 3, 4, 5].

Unfortunately, such a disintermediation can facilitate the spreading of rumors, hoaxes, fake news, and conspiracy theories, that often arouse naive and awkward social responses on different topics, such as health, environment, national security, and politics [6, 7, 8, 9, 10, 11]. In particular, conspiracy theories simplify causation, reduce the complexity of reality, and are formulated in a way

*Corresponding author
Email address: bessi@isi.edu (Alessandro Bessi)

that is able to contain a certain level of uncertainty [12, 13, 14, 15]. However, such an impressive amount of misinformation can influence policy preferences and encourage behaviors strongly divergent from recommended practices.

Since the World Economic Forum listed massive digital misinformation as one of the main threats to our society [16], community-driven [17] and algorithmic-driven [18, 19, 20, 21, 22, 23, 24] solutions have been proposed to counteract the pervasiveness of online misinformation. However, a part of the scientific community is skeptical about the real effectiveness of such solutions. Indeed, the community-driven approach proposed by Facebook – where users can flag false contents to correct the newsfeed algorithm – is controversial, because it raises fears that the free circulation of ideas may be threatened. Moreover, algorithmic-driven approaches may not be effective, since the acceptance of a claim (either substantiated or not) is heavily influenced by social norms and individual cognitive factors [25, 26, 27, 28, 29, 30]. Indeed, recent works point out both the inefficacy of correcting false beliefs and the concrete risk of a backfire effect [31, 32, 33] from the usual and most committed consumers of conspiracy theories. In fact, false beliefs, once adopted by an individual, are rarely corrected [34, 35, 36, 37].

More specifically, both the formation and the revision of beliefs are strongly affected by the communities wherein ideas and facts are debated [38, 39]. Such a phenomenon is emphasized in online social networks, where users process information through a shared system of meaning [40, 41] inside their echo chambers [42, 43, 44, 45], making sense of facts in ways that are often biased toward self-confirmation. Indeed, recent studies show that increasing the exposure of users to unsubstantiated rumors increases their tendency to be credulous [26, 46], and that the content-selective exposure is the primary driver of content diffusion, and generates the formation of echo chambers [42].

In this work, we study the statistical properties of viral misinformation in online social media. In particular, we apply methods of Extreme Value Theory – a branch of statistics dealing with extreme deviations from the median of probability distributions – to analyze a large dataset of posts published by Facebook pages supporting conspiracy theories and myth narratives. By means of an in-depth statistical analysis of the shares distribution and the application of the Peaks Over Threshold (POT) approach, we show that the number of extremely viral posts (e.g. $> 250K$ shares) over time follows a homogeneous Poisson process, and that the interarrival times between such posts are independent and identically distributed, following an exponential distribution. Further, we characterize the uncertainty around the rate parameter of the Poisson process through Bayesian methods. Finally, we are able to derive the predictive posterior probability distribution of the number of posts exceeding a certain threshold of shares over a finite interval of time.

The relevance of our results is not necessarily limited to the field of computational social science coping with misinformation [47, 48, 49]. Indeed, despite the prediction of extremely viral posts and rare events remains an hard task [50, 51], we believe that both our findings and the methodology used herein may be of interest to the broader field of computational social science dealing with forecast-

ing and tracking of viral contents and events [52, 53, 54, 55, 56, 57, 58, 59, 60, 61].

2. Methods

2.1. Ethics Statement

The entire data collection process has been carried out exclusively through the Facebook Graph API, which is publicly available. We used only public available data. The pages from which we downloaded data are public Facebook entities.

2.2. Data Collection

We analyzed 328 US public Facebook pages diffusing conspiratorial beliefs, myth narratives, and controversial information, usually lacking supporting evidence and most often contradictory of the official news. Such a space of investigation is defined with the same approach as in [26, 42], with the support of different Facebook groups very active in monitoring the conspiracy narratives. For each page, we downloaded all the posts (and their respective metadata) in a timespan of 5 years (Jan 1, 2010 to Dec 31, 2014). The dataset is composed by 345, 054 posts. To our knowledge, the dataset is the complete set of conspiracy-like information sources active in the US Facebook scenario up to December 31, 2014.

2.3. Fundamentals of Extreme Value Theory

Extreme value theory (EVT) is a branch of statistics dealing with the extreme deviations from the median of probability distributions. In particular, it aims at assessing the probability of events that are more extreme than any previously observed. Extreme value theory is widely used in many fields of science where power laws play a role in modeling [62], such as structural and geological engineering, finance and risk management, earth sciences, traffic prediction, etc.

In this section, we briefly review some fundamental results of extreme value theory. For extended discussion, proofs, and theorems see [63, 64, 65].

2.3.1. Extreme Value Theory

Suppose $X_1, X_2, \ldots$ are independent and identically distributed (iid) random variables with common cumulative distribution function (cdf) F. Let $M_n = \max\{X_1, \ldots, X_n\}$ denote the maximum of the first n random variables (partial maxima) and let $u(F) = \sup\{x : F(x) < 1\}$ denote the upper endpoint of F. Since

$$\mathbf{Pr}(M_n \leq x) = \mathbf{Pr}(X_1 \leq x, \ldots, X_n \leq x) = F^n(x),$$

M_n converges almost surely to $u(F)$ whether it is finite or infinite. Extreme value theory seeks norming constants $a_n > 0$, $b_n \in \mathbb{R}$, and some nondegenerate distribution function G such that the cdf of the normalized M_n converges to G, i.e.

$$\mathbf{Pr}\left(\frac{M_n - b_n}{a_n} \leq x\right) = F^n(a_n x + b_n) \to^d G(x).$$

If this holds for suitable choices of a_n and b_n, then we say that G is an extreme value distribution function, and F belongs to the maximum domain of attraction of G, i.e. $F \in MDA(G)$. The Extremal Types Theorem characterizes the limit distribution function G as of the type of one of the following three classes:

- Gumbel:
$$\Lambda(x) = \exp\left(\exp(-x)\right), \quad x \in \mathbb{R}$$
- Fréchet:
$$\Phi_\alpha(x) = \begin{cases} 0, & \text{if } x \leq 0 \\ \exp(-x^{-\alpha}), & \text{if } x > 0 \end{cases}$$
- Weibull
$$\Psi_\alpha(x) = \begin{cases} \exp(-(-x)^{\alpha}), & \text{if } x \leq 0 \\ 1, & \text{if } x > 0 \end{cases}$$

for some $\alpha > 0$.

The three extreme value distributions can be represented using the generalized extreme value (GEV) distribution (family). Let

$$H_\xi(x) = \begin{cases} \exp\left(-(1+\xi x)^{-\frac{1}{\xi}}\right), & \text{if } \xi \neq 0 \\ \exp(-\exp(-x)), & \text{if } \xi = 0 \end{cases}$$

where $1 + \xi x > 0$. Then,

- $\xi = \alpha^{-1} > 0 \longleftrightarrow \Phi_\alpha$
- $\xi = -\alpha^{-1} < 0 \longleftrightarrow \Psi_\alpha$
- $\xi = 0 \longleftrightarrow \Lambda$

From a modeling point of view, the three extreme value distributions are very different, especially for what concerns the behavior of the tails, i.e. the part of the distribution more relevant when dealing with extreme events.

Here, we focus on the Fréchet case, Φ_α with $\alpha > 0$. If we consider the tail of $\Phi_\alpha(x)$, a Taylor expansion shows that

$$1 - \Phi_\alpha(x) = 1 - \exp(-x^{-\alpha}) \sim x^{-\alpha}, \quad x \to \infty.$$

Hence, $\Phi_\alpha(x)$ tends to decrease as a power law. Moreover, every distribution function that belongs to MDA(Φ_α) has necessarily and infinite right endpoint, i.e. it is defined for $x \in [0, \infty)$. It follows that all the distribution functions belonging to MDA(Φ_α) are appropriate for modeling phenomena with extremely large maxima.

Consider a random variable X with unknown distribution function G and right endpoint $x_G = \sup\{x \in \mathbb{R} : G(x) < 1\}$. The exceedance distribution function of X above a given threshold t is defined as

$$G_t(x) = \mathbf{Pr}(X \leq x | X > t) = \frac{G(x) - G(t)}{1 - G(t)}, \quad x \geq t.$$

For a large class of distribution functions G and a high threshold $t \to x_G$, G_t can be approximated by a Generalized Pareto Distribution, i.e.

$$G_t = GPD(x; \xi, \beta, t) = \begin{cases} 1 - \left(1 + \xi \frac{x-t}{\beta}\right)^{-\frac{1}{\xi}}, & \text{if } \xi \neq 0 \\ 1 - \exp\left(-\frac{x-t}{\beta}\right), & \text{if } \xi = 0 \end{cases}$$

where $x \geq t$ for $\xi \geq 0$, $t \leq x \leq t - \beta/\xi$ for $\xi < 0$, $t \in \mathbb{R}$, $\xi \in \mathbb{R}$, and $\beta > 0$. The shape parameter, ξ, governs the fatness of the tails, and thus the existence of the moments. The moment of order p of a Generalized Pareto distributed random variable only exists if and only if $\xi < 1/p$.

2.3.2. Extreme Value Analysis

Two approaches exist for practical extreme value analysis. The Block Maxima (BM) approach consists on splitting the observation period into a certain number of non-overlapping periods of equal size – e.g. weeks, months, years – and then considering only the maximal value within each period. Such maximal values follow approximately a Generalized Extreme Value (GEV) distribution.

The Peaks Over Threshold (POT) approach relies on considering only the values exceeding a certain high threshold. The probability distribution of those selected observations is approximately a Generalized Pareto Distribution (GPD).

Both approaches have some limitations. The POT approach picks up all relevant high observations, and thus seems to make better use of the available information. Conversely, the BM approach misses some of these high observations and retains some lower observations. However, there may be reason for using the BM method: the only available information may be block maxima (e.g. daily, weekly, monthly, or yearly maxima) and the BM approach may be preferable when the observations are not exactly iid.

However, if the BM approach may be easier to apply when the block periods appear naturally, some problems arise when this does not happen. In such a case, the choice of the block size for the BM approach may be as difficult as the choice of the threshold for the POT approach.

Figure 1 provides a graphical representation of the two approaches.

3. Results and Discussion

In this paper, we aim at investigating the statistical properties of viral misinformation on Facebook by means of extreme value theory. More specifically, the object of the analysis is the number of times that posts supporting — in

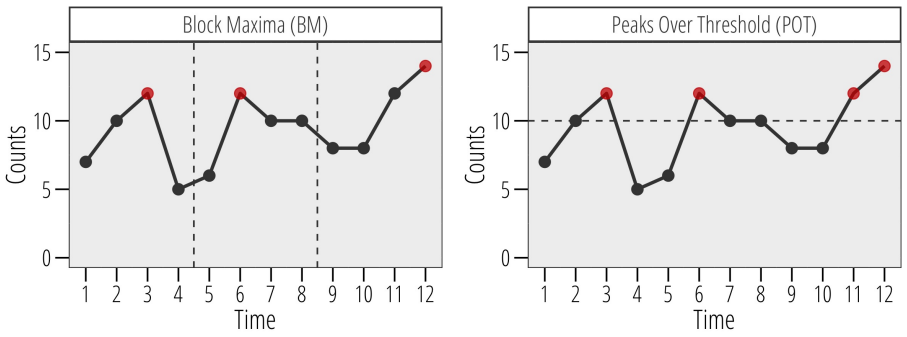

Figure 1: **Block Maxima (BM) vs. Peaks Over Threshold (POT).**

this study, we are assuming that each share of conspiracy posts represents the will to support a given conspiracy narrative — conspiracy theories have been shared, which can be considered as a random variable X following a generic distribution function F with support $[0, \infty)$. Such an infinite right endpoint is justified by the fact that users can share a post how many times they desire.

3.1. Exploratory Data Analysis

Since for each post we know the time of creation, we have a temporally ordered collection of observations. Such a time series is irregularly spaced, in the sense that it is characterized by varying interarrival times between observations. A common approach to analyze irregularly spaced time series consists in transforming the data into equally spaced observations using interpolation methods, and then apply standard methods for equally spaced data. However, such a transformation can introduce a number of significant and hard to quantify bias, especially when the interarrival times between observations are highly irregular. Since in our case the spacing of observations varies from seconds to days, we avoid to transform data. Moreover, despite observations are temporally ordered, it is difficult to assume some kind of time dependence between the number of shares received by posts — i.e. the number of shares received by a post does not affect the number of shares received by following posts. Rather, if we can conceive that some external events (e.g. breaking news, top stories, scandals, etc.) can cause an unusual number of posts in a restricted temporal window (clustering), we can safely assume that the number of shares received by each of those posts is independent and identically distributed (iid).

Figure 2 shows the cumulative number of weekly post (*top panel*) and the number of weekly posts (*bottom panel*). Despite it looks like there is a clear growth trend — which is likely due by the increase of users on Facebook occurring from 2010 to 2014 — and some form of seasonality, we are not able to identify any meaningful seasonality pattern. Indeed, we can assume that the activity of this kind of pages is primarily driven by external events, such as breaking news, top stories, scandals, etc.

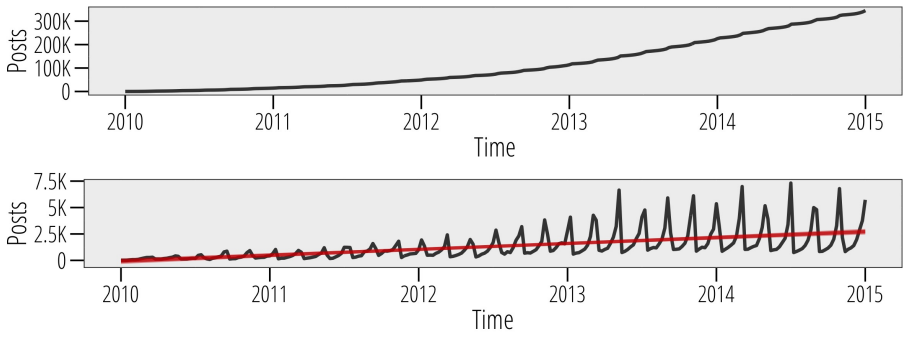

Figure 2: **Weekly posts.** Cumulative number of weekly post (*top panel*) and number of weekly posts (*bottom panel*). The solid red line indicates the fitted linear trend.

However, an increase in the number of posts published by pages in a given temporal window may reflect an increase in the users' excitement and activity. We account for such a possible characteristic of the phenomenon under investigation by rescaling raw data by a factor defined as

$$R_i = \frac{w_i}{\max(w)}, \quad i \in \{1, \ldots, 261\}$$

where w_i represents the number of posts published by pages in week i. Such a rescaling factor inflates the number of shares of posts published in weeks characterized by an overall low activity. Different rescaling strategies have been considered — e.g. rescaling by the mean or the median —, and similar results have been obtained.

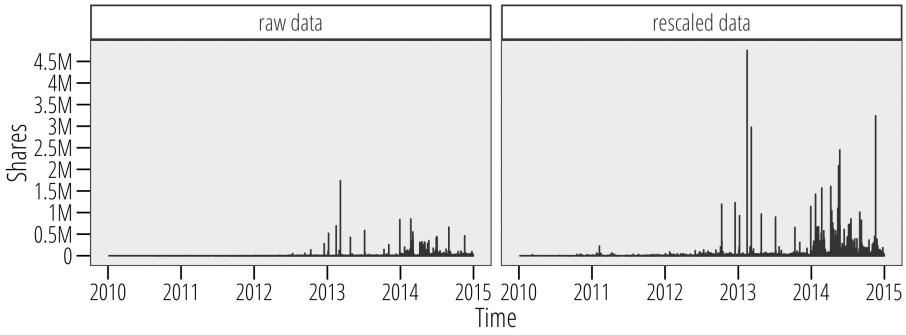

Figure 3: **Shares time series.** Time series of the original number of shares (*raw data*) and the rescaled one (*rescaled data*).

Figure 3 shows the time series of the original number of shares (*raw data*) and the rescaled one (*rescaled data*). The rescaling procedure should have removed or at least reduced possible clustering phenomena that would have led to a violation of the iid assumption. We check the iid assumption by means of the

records plot, a simple and intuitive exploratory tool widely used in extreme value analysis which exploits the fact that successive records for iid data should become more and more rare as time goes by. Since a record x_n for the random variable X occurs if $x_n > \max\{x_1, \ldots, x_{n-1}\}$, it is intuitive that if data are iid it becomes more difficult to exceed all past observations, and thus the number of records should follow a logarithmic pattern [63]. Records plots in Figure 4 show that the iid assumption for raw data is violated, but still valid for rescaled data, where records are distributed around their expected value and within the 95% confidence intervals.

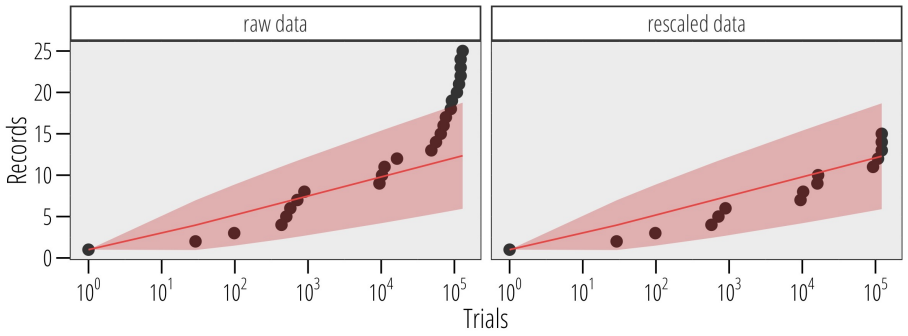

Figure 4: **Records plots.** The iid assumption for raw data is violated, but still valid for rescaled data, where records are distributed around their expected value and within the 95% confidence intervals.

Figure 5 shows the empirical complementary cumulative distribution functions of raw data and rescaled data. The double log scale of the figures highlights the fatness of the right tail in both the two empirical distributions. Beyond removing any form of dependence in the raw data, we observe that the rescaling procedure slightly exacerbates the power law behavior of the tail without influencing the body of the distribution. Thus, rescaled data will be used for successive analysis.

3.2. Statistical Properties of Viral Misinformation

We use the distribution function of rescaled data to characterize the statistical properties of viral misinformation by means of EVT tools. First, we analyze the limit behavior of the Maximum/Sum ratio

$$R_n(p) = \frac{M_n(p)}{S_n(p)}, \qquad n \geq 1, p > 0,$$

where $S_n(p) = \sum_{i=1}^{n}(X_i^p)$ and $M_n(p) = \max(X_i^p)$. The moment of order p of the distribution exists, i.e. $E[X^p] < \infty$, if and only if $R_n(p)$ converges to zero for $n \to \infty$. Conversely, an erratic limit behavior of $R_n(p)$ indicates the infiniteness of the p-th moment of the distribution, i.e. $E[X^p] = \infty$.

Figure 6 shows that only the first moment of the distribution exists, i.e. $E[X] < \infty$, whereas moments of order greater than $p = 2$ are infinite, i.e.

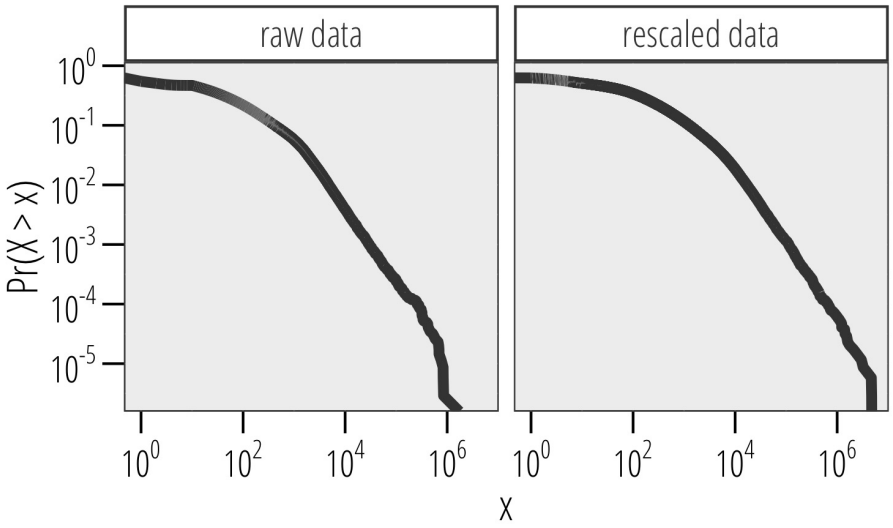

Figure 5: **Empirical Complementary Cumulative Distribution Function.** The rescaling procedure slightly exacerbates the power law behavior of tail without influencing the body of the distribution.

$E[X^p] = \infty$ for $p \geq 2$. Identical results hold for raw data. Our distribution function belongs to the maximum domain of attraction of Fréchet, i.e. $F \in MDA(\Phi_\alpha)$. The existence of the first moment of the distribution function allows us to compute a *reasonable* (in a sample one can compute basically anything, even meaningless quantities) estimate of the conditional tail mean above a given threshold. Indeed, by the law of total expectation

$$E[X|X > t] = \frac{E[X] - \mathbf{Pr}(X \leq t)E[X|X \leq t]}{\mathbf{Pr}(X > t)},$$

where $E[X|X \leq t]$ is finite since bounded from above by t, and the finiteness of $E[X]$ implies that the conditional tail mean, $E[X|X > t]$, is finite.

Such a measure is known in finance as the expected shortfall of a loss distribution, and it let us answer to question such as *"What is the expected number of shares for a post once it has exceeded the* $250K$ *shares threshold?"*. Indeed,

$$E[X|X > t] = \frac{\sum_{i=1}^{N} x_i \mathbb{I}(x_i > t)}{\sum_{i=1}^{N} \mathbb{I}(x_i > t)} = \frac{\sum_{i=1}^{N} x_i \mathbb{I}(x_i > 250K)}{\sum_{i=1}^{N} \mathbb{I}(x_i > 250K)} \approx 467K.$$

Since we showed that the moments of order greater than 1 do not exist, one should prefer the mean absolute deviation over the variance as a measure of dispersion around the conditional tail mean.

Recall that the moment of order p of a Generalized Pareto distributed random variable only exists if and only if $\xi < 1/p$, and thus the shape parameter we are going to estimate can not be smaller than $1/2$. Such an observation has the

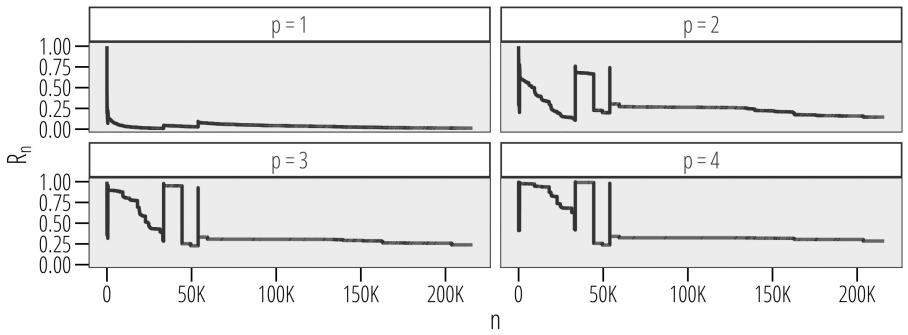

Figure 6: **Maximum/Sum ratio plot.** Only the first moment of the distribution exists, i.e. $E[X] < \infty$, whereas moments of order greater than $p = 2$ are infinite, i.e. $E[X^p] = \infty$ for $p \geq 2$.

main implication that we can safely use the maximum likelihood (ML) approach to estimate ξ, since the ML estimates are consistent only when $\xi > -1/2$.

Since our time series is highly irregularly spaced, with interarrival times ranging between seconds and days, we prefer the Peaks Over Threshold (POT) approach over the Block Maxima (BM) approach to estimate the shape parameter ξ. Indeed, when the block periods used in the BM approach does not appear naturally, the choice of a threshold for the POT approach may be easier.

Before fitting the distribution function to a Generalized Pareto Distribution, we have to identify a feasible threshold. To accomplish such a task, we rely on the Mean Excess Function (MEF). The empirical MEF of a sample of observations $x_1, \ldots, x_n$ is defined as

$$e_n(t) = \frac{\sum_{i=1}^{n}(x_i - t)}{\sum_{i=1}^{n} \mathbb{I}(x_i > t)},$$

that is the ratio between the sum and the number of the exceedances over the threshold t. Figure 7 shows the MEF plot for the rescaled data. We observe that the empirical MEF begins to linearly increase in the threshold at $t \approx 10^4$. Such a behavior characterizes power law distribution functions [66], and thus we choose $t = 10^4$ as threshold.

Given the heavy-tailed behavior of the rescaled data distribution function, the shape parameter ξ is likely to be positive. The left panel of Figure 8 shows the Pickands plot, based on the nonparametric Pickands estimator for ξ, defined as

$$\tilde{\xi}^{(P)}_{\tau,n} = \frac{1}{\log_2} \log \frac{X_{\tau,n} - X_{2\tau,n}}{X_{2\tau,n} - X_{4\tau,n}}, \qquad \tau = 1, \ldots, \lfloor n/4 \rfloor$$

where $X_{\tau,n}$ is the τ-th upper order statistics out of a sample of n observations. The Pickands plot shows a more or less stable behavior of the Pickands estimates for different values of τ, suggesting that the true value of ξ lies in the interval $(0.5, 1)$.

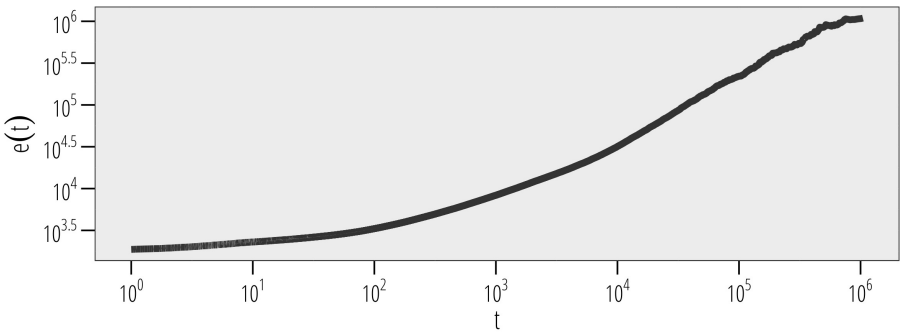

Figure 7: **Empirical Mean Excess Function plot.**

The right panel of Figure 8 shows the Hill plot, based on the nonparametric Hill estimator for ξ, defined as

$$\tilde{\xi}_{\tau,n}^{(H)} = \frac{1}{\tau}\sum_{j=1}^{\tau} \ln(X_{j,n}) - \ln(X_{\tau,n}),$$

where $X_{j,n}$ is the j-th upper order statistics out of a sample of n observations. The Hill plot outperforms the Pickands plot in stability, suggesting a true value of ξ around 0.75.

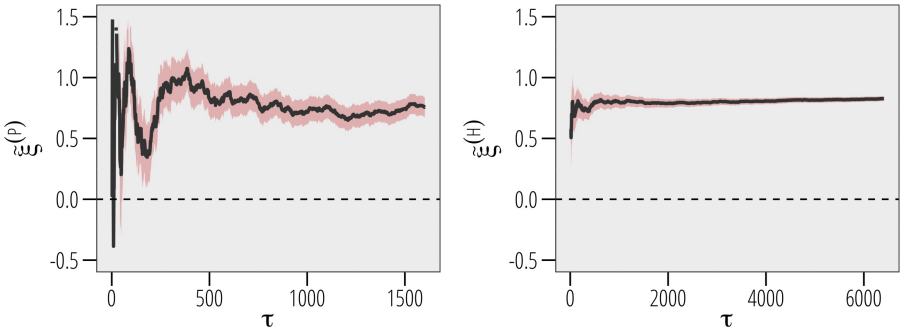

Figure 8: **Pickands and Hill nonparametric estimators for the shape parameter ξ.**

The main implication is that, as already anticipated by the analysis of the Maximum/Sum ratio plot, the true value of ξ is greater than $-1/2$, and thus we can obtain consistent estimates via a maximum likelihood (ML) approach. Table 1 shows ML estimates of ξ and β for different thresholds. We observe a stable value of $\tilde{\xi}^{ML}$ for increasing values of the threshold. We obtain similar results for raw data (i.e. $\tilde{\xi}^{ML} = 0.769(0.0198)$ with $t = 2.5K$).

Table 1: **Maximum Likelihood estimates, standard errors, and number of exceedances for different thresholds.**

threshold	$\tilde{\xi}^{(ML)}$	$\tilde{\beta}^{(ML)}$	# exceedances
$10K$	0.770	8,750	6,408
	(0.0220)	(205)	
$25K$	0.800	19,500	2,153
	(0.0391)	(805)	
$50K$	0.737	43,170	884
	(0.059)	(2,730)	
$100K$	0.746	74,380	399
	(0.0869)	(6,923)	
$150K$	0.726	120,460	223
	(0.120)	(15,500)	

3.3. Frequency of Viral Misinformation

The Peaks Over Threshold (POT) method has two main implications: the exceedances over a high threshold follow a Generalized Pareto Distribution, and the number of excesses over time follows a homogeneous Poisson process. In a homogeneous Poisson process the number of events, $N(\theta)$, in a finite interval of time of length θ follows the Poisson distribution, i.e.

$$\mathbf{Pr}(N(\theta) = n) = \frac{(\lambda\theta)^n}{n!} \exp(-\lambda\theta).$$

Moreover, the interarrival times between events are independent and follow the exponential distribution, i.e.

$$\mathbf{Pr}(\text{interarrival time} > \theta) = \exp(-\lambda\theta).$$

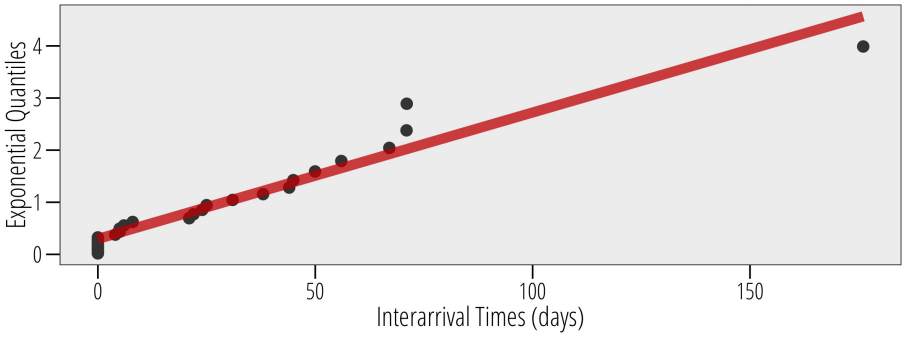

Figure 9: **Exponential Quantiles vs. Interarrival Times.**

Figure 9 shows that the interarrival times of posts shared more than $750K$ times (rescaled data) follow approximately an exponential distribution. Moreover, the autocorrelogram function (ACF) plot — i.e. a plot showing the similarity between observations as a function of the time lag between them [67] — in

Figure 10 shows that the interarrival times between those posts are independent, supporting the i.i.d. hypothesis suggested by the records plot in Figure 4. Similar results approximately hold for raw data when considering posts shared more than $250K$ (the bootstrap test of fit for the Generalized Pareto Distribution [68] gives a p-value equal to 0.2). We conclude that the number of extremely viral posts over time follows a homogeneous Poisson process.

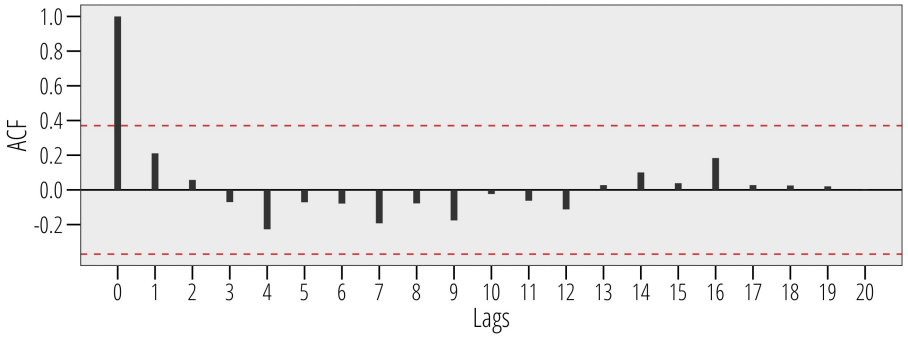

Figure 10: **Autocorrelogram for Interarrival Times.** We find no correlation as a function of the time lag between them.

Such a conclusion allows us to exploit some useful properties of the homogeneous Poisson processes to quantify the frequency of rare viral contents on online social media. Indeed, the expected value of the number of events, $N(\theta)$, in a finite interval of time of length θ is defined as

$$E[N(\theta)] = \lambda\theta,$$

where $\lambda > 0$ is known as the rate parameter of the Poisson process. The reciprocal of such a parameter, i.e. $1/\lambda$, is known as the survival parameter of the exponential distribution followed by the interarrival times between the $N(\theta)$ events. Given a sample $z_1, \ldots, z_n$ of interarrival times, the survival parameter is estimated through the sample mean

$$\frac{1}{\lambda} = \frac{\sum_{i=1}^{n} z_i}{n}.$$

Essentially, we can estimate the survival parameter, $1/\lambda$, of the exponential distribution describing the interarrival times between rare events exceeding a certain threshold, and then use the rate parameter, λ, of the Poisson process to estimate the expected number of events exceeding that threshold in a finite time of length θ.

For instance, the survival parameter of the interarrival times distribution of posts exceeding $250K$ shares (raw data) is $1/\lambda = 18.5$. It follows that $\lambda = 1/18.5 = 0.0541$. Basically, if by means of the survival parameter we can answer to questions such as *"What is the mean waiting time between posts exceeding* $250K$ *shares?"*, through the rate parameter we can answer to questions such as

"What is the expected number of posts exceeding $250K$ *shares in the future* 365 *days?"*. Indeed,

$$E[N(\theta)] = \lambda\theta = 0.0541 \times 365 = 19.8 \approx 20.$$

A convenient way to assess the uncertainty around the rate parameter, λ, consists in using a standard Bayesian probability updating method. Indeed, the conjugate prior distribution for a Poisson distribution is the Gamma distribution, and we can express the prior distribution of λ as

$$\mathbf{Pr}(\lambda) = \text{Gamma}(\alpha, \beta).$$

Since the expected value (mean) of a Gamma distribution is defined as α/β, we may want to choose the hyperparameters, α and β, of the prior distribution $\mathbf{Pr}(\lambda)$ so that

$$\frac{1}{\lambda} = \frac{\beta}{\alpha} = \frac{\sum_{i=1}^{n} z_i}{n},$$

where $z_1, \ldots, z_n$ are the observed interarrival times. Then, the posterior distribution of the rate parameter is defined as

$$\mathbf{Pr}(\lambda|\mathbf{z}) = \text{Gamma}(\alpha + k, \beta + \sum_{i=1}^{k} z_i),$$

where $z_1, \ldots, z_k$ represent k new observed interarrival times. For instance, we may define the prior distribution of the rate parameter of the interarrival times distribution of posts exceeding $250K$ shares (raw data) as

$$\mathbf{Pr}(\lambda) = \text{Gamma}(\alpha, \beta) = \text{Gamma}(n, \sum_{i=1}^{n} z_i) = \text{Gamma}(38, 702),$$

with mean equal to $\alpha/\beta = 38/702 = 0.0541$, and variance equal to $\alpha/\beta^2 = 38/702^2 = 7.71 \times 10^{-5}$. Then, if after 60 days we observe a post exceeding the $250K$ shares threshold, the posterior probability distribution of the rate parameter is

$$\mathbf{Pr}(\lambda|\mathbf{z}) = \text{Gamma}(\alpha + k, \beta + \sum_{i=1}^{k} z_i) = \text{Gamma}(38 + 1, 702 + 60),$$

with mean equal to $(\alpha+k)/(\beta+\sum_{i=1}^{k} z_i) = (38+1)/(702+60) = 0.0512$, and variance equal to $38 + 1/(702 + 60)^2 = 6.72 \times 10^{-5}$. Figure 11 shows both the prior and the posterior distributions of the rate parameter in the aforementioned example.

After such an update, the expected value of the number of posts exceeding the $250K$ threshold in the next 365 days is defined as

$$E[N(\theta)] = E[\lambda|\mathbf{z}]\theta = 0.0512 \times 365 = 18.7 \approx 19,$$

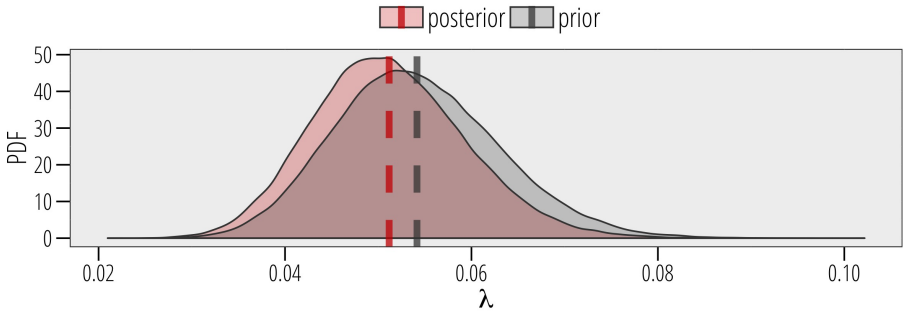

Figure 11: **Prior and Posterior Distributions of the rate parameter.** The grey and red dashed lines indicate, respectively, the mean of the prior and the mean of the posterior distribution of the rate parameter.

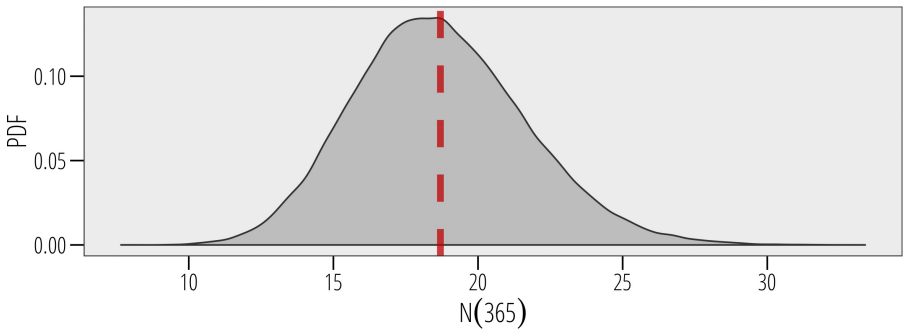

Figure 12: **Posterior predictive probability distribution.** Posterior predictive probability distribution of the number of posts exceeding the $250K$ shares threshold in the next 365 days. The red dashed line indicates the mean, i.e. 18.7.

and the mean waiting time between posts exceeding the $250K$ shares is $1/0.0512 = 19.5$ days. Moreover, the full probability assessment of the uncertainty around the rate parameter, λ, allows us to express the predictive posterior probability distribution of the number of posts exceeding a certain threshold over a finite interval of time

$$\mathbf{Pr}(N(\theta)|\lambda) = \mathbf{Pr}(\lambda|\mathbf{z})\theta.$$

Figure 12 shows the posterior predictive probability distribution function of the number of posts exceeding the $250K$ shares threshold (raw data) in the finite interval time of length 365 days.

3.4. Concluding Remarks

In this paper, we study the statistical properties of viral misinformation in online social media. In particular, we focus our attention on Facebook posts

spreading false news, hoaxes and unsubstantiated claims. By means of an Extreme Value Theory approach, we show that the number of extremely viral posts over time follows a homogeneous Poisson process, and that the interarrival times between such posts are independent and identically distributed, following an exponential distribution. Moreover, we characterize the uncertainty around the rate parameter of the Poisson process through Bayesian methods. Finally, we are able to derive the predictive posterior probability distribution of the number of posts exceeding a certain threshold of shares over a finite interval of time.

The relevance of our results is not necessarily limited to the field of computational social science coping with misinformation. Despite the prediction of extremely viral posts — and, more generally, rare events — remains an hard task, we believe that both our findings and the methodology introduced in this paper may be of interest to the broader field of computational social science dealing with forecasting and tracking of viral contents and events — e.g. cybersecurity attacks, terrorist attacks, etc.

Acknowledgements

Special thanks to Geoff Hall and Skepti Forum for providing fundamental support in defining the atlas of Facebook pages disseminating conspiracy theories and myth narratives.

References

References

[1] J. Brown, A. J. Broderick, N. Lee, Word of mouth communication within online communities: Conceptualizing the online social network, Journal of interactive marketing 21 (3) (2007) 2–20.

[2] R. Kahn, D. Kellner, New media and internet activism: From the'battle of seattle'to blogging., New media & society 6 (1) (2004) 87–95.

[3] W. Quattrociocchi, R. Conte, E. Lodi, Opinions manipulation: Media, power and gossip, Advances in Complex Systems 14 (04) (2011) 567–586.

[4] W. Quattrociocchi, G. Caldarelli, A. Scala, Opinion dynamics on interacting networks: media competition and social influence, Scientific reports 4.

[5] R. Kumar, M. Mahdian, M. McGlohon, Dynamics of conversations, in: Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2010, pp. 553–562.

[6] Ebola Lessons: How Social Media Gets Infected, `http://www.informationweek.com/software/social/-ebola-lessons-how-social-media-gets-infected/a/d-id/1307061` (March 2014).

[7] The Ebola Conspiracy Theories, http://www.nytimes.com/2014/10/19/sunday-review/the-ebola-conspiracy-theories.html (March 2014).

[8] The inevitable rise of Ebola conspiracy theories, http://www.washingtonpost.com/blogs/wonkblog/wp/2014/10/13/the-inevitable-rise-of-ebola-conspiracy-theories/ (March 2014).

[9] Remember jade helm 15, the controversial military exercise? its over (March 2014) [cited 29.10.2014].
URL https://www.washingtonpost.com/news/checkpoint/wp/2015/09/14/remember-jade-helm-15-the-controversial-military-exercise-its-over/

[10] 5 myths surrounding vaccines — and the reality (February 2015).
URL http://edition.cnn.com/2015/02/04/us/5-vaccine-myths/

[11] Trumps outrageous claim that thousands of new jersey muslims celebrated the 9/11 attacks (November 2015).
URL https://www.washingtonpost.com/news/fact-checker/wp/2015/11/22/donald-trumps-outrageous-claim-that-thousands-of-new-jersey-muslims-celebrated-th

[12] C. R. Sunstein, A. Vermeule, Conspiracy theories: Causes and cures*, Journal of Political Philosophy 17 (2) (2009) 202–227.

[13] J. Byford, Conspiracy theories: a critical introduction, Palgrave Macmillan, 2011.

[14] G. A. Fine, V. Campion-Vincent, C. Heath, Rumor mills: The social impact of rumor and legend, Transaction Publishers, 2005.

[15] M. A. Hogg, D. L. Blaylock, Extremism and the Psychology of Uncertainty, Vol. 8, John Wiley & Sons, 2011.

[16] L. Howell, Digital wildfires in a hyperconnected world, WEF Report.

[17] News feed fyi: Showing fewer hoaxes (January 2015).
URL http://newsroom.fb.com/news/2015/01/news-feed-fyi-showing-fewer-hoaxes/

[18] V. Qazvinian, E. Rosengren, D. R. Radev, Q. Mei, Rumor has it: Identifying misinformation in microblogs, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, 2011, pp. 1589–1599.

[19] G. L. Ciampaglia, P. Shiralkar, L. M. Rocha, J. Bollen, F. Menczer, A. Flammini, Computational fact checking from knowledge networks, PloS one 10 (6) (2015) e0128193.

[20] P. Resnick, S. Carton, S. Park, Y. Shen, N. Zeffer, Rumorlens: A system for analyzing the impact of rumors and corrections in social media, in: Proc. Computational Journalism Conference, 2014.

[21] A. Gupta, P. Kumaraguru, C. Castillo, P. Meier, Tweetcred: Real-time credibility assessment of content on twitter, in: Social Informatics, Springer, 2014, pp. 228–243.

[22] A. A. AlMansour, L. Brankovic, C. S. Iliopoulos, A model for recalibrating credibility in different contexts and languages-a twitter case study, International Journal of Digital Information and Wireless Communications (IJDIWC) 4 (1) (2014) 53–62.

[23] J. Ratkiewicz, M. Conover, M. Meiss, B. Gonçalves, A. Flammini, F. Menczer, Detecting and tracking political abuse in social media., in: ICWSM, 2011.

[24] X. L. Dong, E. Gabrilovich, K. Murphy, V. Dang, W. Horn, C. Lugaresi, S. Sun, W. Zhang, Knowledge-based trust: Estimating the trustworthiness of web sources, Proceedings of the VLDB Endowment 8 (9) (2015) 938–949.

[25] D. Mocanu, L. Rossi, Q. Zhang, M. Karsai, W. Quattrociocchi, Collective attention in the age of (mis) information, Computers in Human Behavior 51 (2015) 1198–1204.

[26] A. Bessi, M. Coletto, G. A. Davidescu, A. Scala, G. Caldarelli, W. Quattrociocchi, Science vs conspiracy: Collective narratives in the age of misinformation, PloS one 10 (2) (2015) e0118093.

[27] B. Nyhan, J. Reifler, S. Richey, G. L. Freed, Effective messages in vaccine promotion: a randomized trial, Pediatrics 133 (4) (2014) e835–e842.

[28] M. A. Javarone, Social influences in opinion dynamics: the role of conformity, Physica A: Statistical Mechanics and its Applications 414 (2014) 19–30.

[29] M. A. Javarone, G. Armano, Perception of similarity: a model for social network dynamics, Journal of Physics A: Mathematical and Theoretical 46 (45) (2013) 455102.

[30] M. A. Javarone, Network strategies in election campaigns, Journal of Statistical Mechanics: Theory and Experiment 2014 (8) (2014) P08013.

[31] A. Bessi, G. Caldarelli, M. Del Vicario, A. Scala, W. Quattrociocchi, Social determinants of content selection in the age of (mis) information, in: Social Informatics, Springer, 2014, pp. 259–268.

[32] F. Zollo, A. Bessi, M. Del Vicario, A. Scala, G. Caldarelli, L. Shekhtman, S. Havlin, W. Quattrociocchi, Debunking in a world of tribes, arXiv preprint arXiv:1510.04267.

[33] R. Garrett, E. Nisbet, E. Lynch, Undermining the corrective effects of media-based political fact checking? the role of contextual cues and naive theory, Journal of Communication.

[34] R. K. Garrett, B. E. Weeks, The promise and peril of real-time corrections to political misperceptions, in: Proceedings of the 2013 conference on Computer supported cooperative work, ACM, 2013, pp. 1047–1058.

[35] M. L. Meade, H. L. Roediger, Explorations in the social contagion of memory, Memory & cognition 30 (7) (2002) 995–1009.

[36] A. Koriat, M. Goldsmith, A. Pansky, Toward a psychology of memory accuracy, Annual review of psychology 51 (1) (2000) 481–537.

[37] M. S. Ayers, L. M. Reder, A theoretical review of the misinformation effect: Predictions from an activation-based memory model, Psychonomic Bulletin & Review 5 (1) (1998) 1–21.

[38] B. Zhu, C. Chen, E. F. Loftus, C. Lin, Q. He, C. Chen, H. Li, R. K. Moyzis, J. Lessard, Q. Dong, Individual differences in false memory from misinformation: Personality characteristics and their interactions with cognitive abilities, Personality and Individual Differences 48 (8) (2010) 889–894.

[39] S. J. Frenda, R. M. Nichols, E. F. Loftus, Current issues and advances in misinformation research, Current Directions in Psychological Science 20 (1) (2011) 20–23.

[40] A. Bessi, F. Zollo, M. Del Vicario, A. Scala, G. Caldarelli, W. Quattrociocchi, Trend of narratives in the age of misinformation, PloS one 10 (8) (2015) e0134641.

[41] F. Zollo, P. K. Novak, M. Del Vicario, A. Bessi, I. Mozetič, A. Scala, G. Caldarelli, W. Quattrociocchi, Emotional dynamics in the age of misinformation, PloS one 10 (9) (2015) e0138740.

[42] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H. E. Stanley, W. Quattrociocchi, The spreading of misinformation online, Proceedings of the National Academy of Sciences 113 (3) (2016) 554–559.

[43] A. Bessi, F. Petroni, M. Del Vicario, F. Zollo, A. Anagnostopoulos, A. Scala, G. Caldarelli, W. Quattrociocchi, Viral misinformation: The role of homophily and polarization, in: Proceedings of the 24th International Conference on World Wide Web Companion, International World Wide Web Conferences Steering Committee, 2015, pp. 355–356.

[44] A. Bessi, F. Zollo, M. Del Vicario, M. Puliga, A. Scala, G. Caldarelli, B. Uzzi, W. Quattrociocchi, Users polarization on facebook and youtube, arXiv preprint arXiv:1604.02705.

[45] A. Bessi, Personality traits and echo chambers on facebook, Computers in Human Behavior 65 (2016) 319–324.

[46] A. Bessi, A. Scala, L. Rossi, Q. Zhang, W. Quattrociocchi, The economy of attention in the age of (mis) information, Journal of Trust Management 1 (1) (2014) 1–13.

[47] C. Shao, G. L. Ciampaglia, A. Flammini, F. Menczer, Hoaxy: A platform for tracking online misinformation, in: Proceedings of the 25th International Conference Companion on World Wide Web, International World Wide Web Conferences Steering Committee, 2016, pp. 745–750.

[48] D. R. Grimes, On the viability of conspiratorial beliefs, PloS one 11 (1) (2016) e0147905.

[49] A. Zubiaga, M. Liakata, R. Procter, K. Bontcheva, P. Tolmie, Towards detecting rumours in social media, arXiv preprint arXiv:1504.04712.

[50] M. J. Salganik, P. S. Dodds, D. J. Watts, Experimental study of inequality and unpredictability in an artificial cultural market, science 311 (5762) (2006) 854–856.

[51] D. J. Watts, Everything is obvious: How common sense fails us, Crown Pub, 2011.

[52] J. Cheng, L. Adamic, P. A. Dow, J. M. Kleinberg, J. Leskovec, Can cascades be predicted?, in: Proceedings of the 23rd international conference on World wide web, ACM, 2014, pp. 925–936.

[53] A. Friggeri, L. A. Adamic, D. Eckles, J. Cheng, Rumor cascades., in: ICWSM, 2014.

[54] J. Staiano, D. Albanese, et al., Exploring image virality in google plus, in: Social Computing (SocialCom), 2013 International Conference on, IEEE, 2013, pp. 671–678.

[55] T.-A. Hoang, E.-P. Lim, Virality and susceptibility in information diffusions., in: ICWSM, 2012.

[56] L. Hong, O. Dan, B. D. Davison, Predicting popular messages in twitter, in: Proceedings of the 20th international conference companion on World wide web, ACM, 2011, pp. 57–58.

[57] M. Jenders, G. Kasneci, F. Naumann, Analyzing and predicting viral tweets, in: Proceedings of the 22nd international conference on World Wide Web companion, International World Wide Web Conferences Steering Committee, 2013, pp. 657–664.

[58] J. Yang, S. Counts, Predicting the speed, scale, and range of information diffusion in twitter., ICWSM 10 (2010) 355–358.

[59] M. Coscia, Average is boring: How similarity kills a meme's success, Scientific reports 4.

[60] L. Weng, F. Menczer, Y.-Y. Ahn, Virality prediction and community structure in social networks, Scientific reports 3.

[61] J. Zhou, H. Pei, H. Wu, Early warning of human crowds based on query data from baidu map: Analysis based on shanghai stampede, arXiv preprint arXiv:1603.06780.

[62] P. Cirillo, N. N. Taleb, On the statistical properties and tail risk of violent conflicts, Available at SSRN 2675355.

[63] E. Gumbel, Statistics of extremes. 1958, Columbia Univ. press, New York.

[64] S. Coles, J. Bawa, L. Trenner, P. Dorazio, An introduction to statistical modeling of extreme values, Vol. 208, Springer, 2001.

[65] P. Embrechts, C. Klüppelberg, T. Mikosch, Modelling extremal events: for insurance and finance, Vol. 33, Springer Science & Business Media, 2013.

[66] P. Cirillo, Are your data really pareto distributed?, Physica A: Statistical Mechanics and its Applications 392 (23) (2013) 5947–5962.

[67] G. E. Box, G. M. Jenkins, G. C. Reinsel, G. M. Ljung, Time series analysis: forecasting and control, John Wiley & Sons, 2015.

[68] J. A. Villaseñor-Alva, E. González-Estrada, A bootstrap goodness of fit test for the generalized pareto distribution, Computational Statistics & Data Analysis 53 (11) (2009) 3835–3841.